

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
22 February 2001 (22.02.2001)

PCT

(10) International Publication Number
WO 01/13585 A1

(51) International Patent Classification⁷: H04L 12/46

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.

(21) International Application Number: PCT/US00/22470

(22) International Filing Date: 16 August 2000 (16.08.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/375,255 16 August 1999 (16.08.1999) US
09/512,127 24 February 2000 (24.02.2000) US

(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).


(71) Applicant: INTERNAP NETWORK SERVICES, INC.
[US/US]; Two Union Square, 601 Union Street, Suite 1000,
Seattle, WA 98101-4064 (US).

Published:
— With international search report.

(74) Agent: DEHAAN, May, Lin; Lariviere, Grubman &
Payne, LLP, P.O. Box 3140, Monterey, CA 93942 (US).

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: PRIVATE NETWORK ACCESS POINT ROUTER FOR INTERCONNECTING AMONG INTERNET ROUTE PROVIDERS



WO 01/13585 A1

(57) Abstract: An improvement to the Private Network Access Point (PNAP) packet switched network described in U.S. Patent No. 6,009,081, where two customers (1, 2) connected to the same PNAP (20) will exchange traffic through the PNAP (20) without transiting over the backbones of the Internet (22). In addition, a multi-homed customer connected to the PNAP is provided with access to the PNAP optimized routing table so that the customer will also have the ability to know the best route for a particular destination. In this way, if a multi-homed customer connected to the PNAP is directly connected to a particular NSP to which a destination is also connected, the PNAP customer can use the PNAP information regarding the NSP to send the information to the destination through that commonly connected to NSP in the most direct fashion.

PRIVATE NETWORK ACCESS POINT ROUTER FOR INTERCONNECTING AMONG INTERNET ROUTE PROVIDERS

5

CROSS-REFERENCE TO RELATED APPLICATIONS

This application claims priority to US patent application serial number 09/512,127, filed February 24, 2000. Which is a continuation-in-part of copending application serial number 09/375,255 filed on August 16, 1999, incorporated herein by reference, which is a continuation of serial number 08/922,954 filed on September 3, 1997, now U.S. Patent No. 10 6,009,081, incorporated herein by reference. Priority is claimed to 09/512,127 and 09/375,255.

BACKGROUND OF THE INVENTION

Field of the Invention

15 This invention pertains generally to routing information packets in a network involving a plurality of traffic carrying networks, and more particularly to an improvement in routing described in U.S. Patent 6,009,081.

Description of the Background Art

The present invention is an improvement on the invention of improvement in routing 20 described in U.S. Patent 6,009,081, and assigned to the assignee hereof. Additional background information can be found the aforesaid patent, as well as in the book entitled "Internet Routing Architectures" by Bassam Halabi, New Riders Publishing, 1997, which is hereby incorporated herein by reference.

As indicated in U.S. Patent No. 6,009,081, column 6, lines 62-66, a PNAP or "Private 25 Network Access Point" can be thought of as being made up of two halves. One half connects to customers. The other half connects to NSPs or "National Service Providers".

The Internet is a network of networks. A PNAP contains an ASimilater that determines the Internet interconnection matrix. ASimilater servers residing within the PNAP collect and collate the routing data received from Network Service Providers (NSPs) to build a database 30 of how the Internet is interconnected. The database shows the NSPs connected to the PNAP are interconnected as well as how they are connected to their customers. The PNAP receives

each NSP's perspective of the Global Routing Table which, when collated, includes identical routes from multiple NSPs, and that distillation of the sum of each NSP's view of the Global Routing Table is used to direct traffic from the customer to the destination over the optimal path via another PNAP customer if available or, otherwise, one of the NSP's connected to the
5 PNAP.

BRIEF SUMMARY OF THE INVENTION

In accordance with an aspect of the invention, if two customers who are connected to the same PNAP wish to communicate with each other, traffic will be exchanged between those customers through the PNAP without ever transiting over the backbones of the NSPs.

10 According to another aspect of the invention, a multi-homed customer connected to the PNAP is provided with access to the PNAP's optimized version of the Global Routing Table so that the customer will also have the ability to know the best route for a particular destination.

According to a still further aspect of the invention, if a multi-homed customer
15 connected to the PNAP is directly connected to a particular NSP to which a destination is also connected, the PNAP customer can, based on information provided by the PNAP, send the information to the destination through that commonly connected NSP.

According to another aspect of the invention, provision is made for the routing of traffic for customers who are multi-homed to multiple PNAPs in addition to one or more of the
20 commonly connected NSPs.

A further aspect of the invention provides for routing traffic for customers who are not massively multi-homed, but are connected to more than one PNAP.

Further objects and advantages of the invention will be brought out in the following portions of the specification, wherein the detailed description is for the purpose of fully
25 disclosing preferred embodiments of the invention without placing limitations thereon.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention will be more fully understood by reference to the following drawings which are for illustrative purposes only:

FIG. 1 is a schematic diagram showing two PNAPs with multi-homed customers
30 according to an embodiment of the invention.

FIG. 2 and FIG. 3 are flow charts showing a method of the invention for causing traffic

between two customers of the same PNAP to be exchanged through the PNAP without transiting over the Internet.

FIG. 4 is a schematic diagram of a customer multi-homed to a PNAP and a plurality of NSPs in accordance with the invention.

5

DESCRIPTION OF THE INVENTION

Referring more specifically to the drawings, for illustrative purposes the present invention will now generally be described in connection with the system configuration, setup and operational methodology shown in FIG. 1 through FIG. 4. It will be appreciated that the 10 system may vary as to configuration, and that the method may vary as to the specific steps and sequence, without departing from the basic inventive concepts disclosed herein.

Referring first to FIG. 1, in accordance with an embodiment of the invention a first PNAP 20 and second PNAP 20a are both shown as a circle with a vertical dashed line 21 dividing it in half. While more customers would typically be connected to the PNAPs 20, 20a, 15 the left half of the PNAPs 20, 20a are shown connected to two customers 1, 2 as an example to simplify the discussion. Furthermore, while two PNAPs 20, 20a are shown, there could be one or any other number of PNAPs. While customers 1, 2 are both shown connected to two PNAPs, a customer may be connected only to one PNAP or any other number of PNAPs. Note also in this regard, that the two customers are considered to be "multi-homed" because 20 they are connected to more than one PNAP. In addition, customers 1, 2 each have a second link that is connected to the Internet 22. This is also considered to be a "multi-homed" configuration. It will be appreciated, however, that it is not necessary for either customer 1 or customer 2 to be multi-homed to employ all of the inventive attributes described herein. Also, while PNAP 20 will be referred to herein for simplicity, the discussion is applicable to 25 PNAP 20a as well.

In the configuration shown, the right half of the PNAP 20 is connected to a plurality of NSPs A, B, C, D, ... N which, in turn, form the Internet 22 to which Internet users, such as destinations 3, 4 are also connected. Note that the NSPs A-N do not exchange traffic among themselves through the PNAP 20. Traffic exchanges between NSPs A-N takes place at public 30 or private peering points (not shown).

The customers 1, 2 typically route their traffic through the PNAP 20 from the left half

to the right half. The PNAP 20 then selects the path from the customers 1, 2 to the destinations 3, 4.

- From U.S. Patent No. 6,009,081, it will be understood that the PNAP 20 contains an ASimilator that determines how everyone on the Internet 22 is connected to everyone else.
- 5 Hereinafter, the term "ASimilator" will be used synonymously with the term "A~~s~~similator" in that patent. It will also be understood that the Border Gateway Protocol, version 4 protocol (BGP4) used therein encompasses the concept of a "Global Routing Table" which may be defined as the list of all routes visible to each provider, both of its customers as well as its peers and their customers, of everyone to which they are connected. Briefly, an ASimilator
10 server inside the PNAP 20 receives a data "dump" of the Global Routing Table from each of its NSPs A-N, and collates the data together to build a database of how the Internet 22 is interconnected. The database shows how all of the NSPs A-N are connected together as well as connections to their customers. Once the ASimilator has populated this database, it uses the forward path or reverse path algorithm defined in U.S. Patent No. 6,009,081 to determine
15 which routes are NSP A's customers, which are NSP B's customers, and so on, for all of the NSPs. In effect, the ASimilator "mines" this database. To summarize:
1. The ASimilator takes a dump of the Global Routing Table from each NSP A-N.
 2. The ASimilator collates the data from each NSP's perspective of the Global Routing Table.
 - 20 3. The ASimilator builds a summed Global Routing Table database of the Internet 22's interconnection matrix.
 4. The ASimilator determines which routes are NSP A's customers and so on for all customers and for all other NSPs B-N. As a clarification, note that each NSP is also sending routes of all other NSPs to which it is connected.
 - 25 The routing table inside the PNAP 20 also maps a plurality of routes from customer 1 to customer 2 that go through the NSPs A-N.

In accordance with the present invention, if neither customer 1 nor customer 2 is multi-homed and those customers wish to communicate with each other, traffic will be exchanged between those customers through the PNAP 20 without ever transiting over the backbones of
30 the NSPs A-N. In the case of sending information from customer 1 to customer 2, the routing table inside the PNAP 20 would list the direct connection from customer 1 to customer 2

through the left half of the PNAP 20 over the dotted path 25 as the optimum route. This means that communications between customer 1 and customer 2 who are connected to the PNAP 20 would always use the dotted path 25 as the preferred path unless a failure or flaw prevents that path from being used, in which case traffic between those customers would be exchanged 5 through the Internet.

Accordingly, data packets would typically flow from customer 1 to the PNAP 20 and directly to customer 2 without traversing any of the NSPs A-N that comprise the Internet 22. This is illustrated in FIG. 2 and FIG. 3, which are method flow charts. In FIG. 2, the method begins at block 30 and proceeds to block 31 which is the step of causing the router within the 10 PNAP to list the direct route through the PNAP as one of its routes between two customers connected to the PNAP. The step of the next block 32 is causing the level of preference for the direct path to be higher than for any other routes between the two customers. The step of the next block 33 is causing the router protocol to select the direct route as being the best path between the two customers. Finally, the last block 34 of FIG. 2 is "end". Similarly for FIG. 15 3, the method begins at block 36. The first step in block 37 is causing the customer router to forward a packet from customer 1 to customer 2 over the PNAP link. The next step in block 38 is causing the PNAP router to forward the packet from customer 1 over the direct PNAP path to customer 2 without transiting a service provider backbone. Finally the last block 40 of FIG. 3 is "end".

20 The potential for unacceptable path latency is reduced by this direct connection between customer 1 and customer 2. Path latency can, for example, result from delay between the time when a device receives a frame and the time that frame is forwarded out the destination port, or the delay caused by a shift to a more circuitous path due to an outage.

25 With regard to exchanging information between, for example, customers 1, 2 and destinations 3, 4, usually there will be more than one route from the customers 1, 2 to the destinations 3, 4. Therefore, the routers within the PNAP are used to forward packet traffic through the Internet 22 in an optimized fashion. The routers build routing tables that contain their distillation of the summed Global Routing Table resulting in the best paths to all the 30 destinations from the PNAP's perspective. They both advertise and receive route information to and from other routers. The routers keep track of next hop information that enables a data

packet to reach its destination. A router that does not have a direct physical connection to the destination checks its routing table and forwards the packet to its next hop; that is, a router that it is directly connected to and is closer to that destination. This process repeats until the traffic reaches its destination.

5 In a multi-homed configuration as shown in FIG. 1, if customer 1 wishes to send a packet to destination 3, it will see a link 23 to the Internet 22 and a link 24 to the PNAP 20. As part of the BGP4 protocol, customer 1 inherently has complete control over the outbound routing of its traffic communications in this configuration. As such, said customer may set the BGP4 local preference on the routes received by its router to destination 3 in order to cause
10 it to prefer a particular link. For example, if destination 3 is connected to NSP D, customer 1 may set the preferences within its router to prefer link 24 based on link 24 being the optimum link. Otherwise, link 23 to NSP D may be the preferred link. However, in the event that a fault or failure appears on the preferred link, diversity considerations will cause the other link to be used instead.

15 In order for the customer to be able to set the preferences within its router to cause it to prefer a particular link, the customer needs routing information to know which path is optimum. Therefore, in a multi-homed configuration with the PNAP and another provider, the customer is given access to ASimilater data over its BGP feed to the PNAP. This is done so that the PNAP customer can effectively use both their PNAP and their other NSP pipe.
20 Without the additional ASimilator data in the form of BGP communities on the customer's BGP feed from the PNAP, they are left with attempting to push traffic over the PNAP and provider pipes in a sub-optimal fashion. Again, it may be preferred for a customer to use its pipe to NSP D for communicating with destinations that are connected to NSP D and to use the PNAP (and its external connections to NSPs A-N) for all other destinations. The optimized and
25 distilled Global Routing Table would be sent to the PNAP customer. In this example, the BGP4 attribute known as the "community" would be used to tag NSP C customer routes as determined by ASimilater with the PNAP NSP C customer community. Since the customer has complete control over outbound traffic, the customer can set the local preference in its router to tag a particular route of multiple identical routes from multiple sources as the
30 preferred route. The higher the local preference, the more preferred the route. For example, on the inbound policy applied to the routes received from the PNAP, any routes tagged with

the PNAP's community for NSP D could have their local preference set to 50 and every other route (not tagged) set to 150. On the BGP feed from NSP D, the customer could leave all routes at local preference 100 which is the default. This allows the customer to optimize their routing so that the direct pipe to NSP D is used for destinations on NSP D and the PNAP 20
5 is used for other destinations, thus providing effective and optimized use of both the customer's PNAP and NSP pipes based on the ASimilater information related to said customer over the PNAP BGP feed.

On the other hand, when the preferred link is over the PNAP 20 (e.g., when destination 3 is not a customer of NSP D to which customer 1 is also connected), the data packet is
10 transmitted from customer 1 over link 24 to the left half of the PNAP 20. The PNAP routing infrastructure within the PNAP 20 will have determined a plurality of paths to destination 3. These different paths to the same destination are listed in a routing table along with a parameter indicating the degree of preference attached to each route of a set of the different paths. By manipulating the local preference component of the route selection process of the BGP4
15 protocol, the PNAP 20 picks the best path for the traffic to traverse to reach destination 3. The data packet leaves the right side of the PNAP 20 via the selected one of the NSPs A-N, follows the selected best path through the Internet 22, and reaches destination 3.

Therefore, in accordance with the present invention, two customers connected to the
20 same PNAP 20 see the PNAP 20 as the best path, and exchange traffic with each other through the PNAP 20 without ever going out over the backbones of the NSPs A-N. Or, if a PNAP customer is directly connected to a particular NSP to which a destination is also connected, the PNAP customer can utilize that NSP connection to send the traffic to the destination based on the ASimilater information received over the BGP peering with the PNAP.
25

In the case of sending information from customer 1 to customer 2, the routing table inside the PNAP 20 would list the direct connection from customer 1 to customer 2 through the left half of the PNAP 20 over the dotted path 25 as the optimum route. This means that communications between customer 1 and customer 2 who are connected to the PNAP 20 should always use the dotted path 25 as the preferred path unless a failure or flaw prevents that path
30 from being used.

Thus far we have described what will be referred to as "generic" Diversity +. When

a PNAP customer is multi-homed to more than one PNAP and one NSP, routing outbound traffic become increasingly complex. By way of additional background, the invention described in U.S. Patent No. 6,009,081 subscribes to the model of symmetrical routing of traffic. This method allows us to bypass the public NAPs for approximately ninety per cent 5 of the traffic flowing in and out of our PNAPs with the associated benefit of much higher performance than is normal experienced in today's Internet.

The way we accomplish this symmetrical when optimal routing of traffic is by use of our routing technology called ASimilater. Each PNAP has it's own BGP AS and is completely distinct from the routing perspective of the other PNAPs with no private backbone connecting 10 the PNAPs.

Each PNAP is, however, connected to the same fabric of NSPs as all other PNAPs. The levels of bandwidth to a PNAP may be larger or smaller depending on it's location but the fabric is the same. With that in mind, let us examine the routing of PNAP-SFJ as an example.

15 First, assume that each PNAP is connected to the same fabric of NSPs as all other PNAPs. Generally speaking, routing of traffic inbound from an NSP over the pipe to said NSP is easy. All of these NSPs attach a higher local preference to the routes heard from their customers over those same routes heard from their peers. Routing outbound traffic in a massively multi-homed network is much more difficult. Faced with such a multiplicity of 20 links, the question of how to route traffic in a tightly controlled fashion is one of great importance in attaining the highest performance.

Note that we do not peer with the NSPs that we connect to, but are, full transit customers of each one. This allows us to receive each NSPs 25 perspective on the global routing table. ASimilater collates that data together and builds an interconnection matrix of the entire Internet. With that information, ASimilater can then route traffic optimally from each PNAP.

An additional function of ASimilater is to control the inter-PNAP routing. We optimize the connectivity between the PNAPs as well since we can use any of the NSPs connecting the PNAPs to route traffic between them. This allows us to choose the fastest NSP between any 30 two PNAPs, and thus allows us to offer the optimal path between our customers and the Internet.

In the case of Diversity +, we offer our customers access to ASimilater data over their BGP feed to the PNAP 20 by use of the BGP community attribute. In other words, if a customer is connected to NSP C and a PNAP, we can offer our customer all of NSP C and NSP C's customers routes tagged with a specific community InterNAP community, in this case 5 6993:XXX.

That information allows our customer to route traffic destined to NSP C and NSP C customers over the NSP C link and all other traffic routed over the PNAP connection. This allows a customer to enjoy the same performance gains of symmetrical routing of traffic as PNAP even over a pipe not connected to the PNAP 20.

10 Referring also to FIG. 4, in the customer 5 topology there is a connection to NSP A, a connection to NSP B and a connection to InterNAP (PNAP-SFJ). In this topology, we recommend the following configuration:

- (a) NSP A customer routes over the NSP A link.
- (b) NSP B customer routes over the NSP B link.
- 15 (c) All others over the PNAP link.

In order for this to occur, we send the customer NSP A and NSP B routes tagged with the following communities:

NSP A: 6993:NSP A

NSP B: 6993:NSP B

20 For clarity, let's create the table of the local pref values to use in our IBGP.

Table 1

| | | NSP B | NSP A | PNAP |
|----|-------|-------|-------|------|
| 25 | NSP B | 80 | 45 | 75 |
| | NSP A | 40 | 90 | 75 |
| | PNAP | 40 | 45 | 150 |

Setting the fall-through local pref values to half of the primary assists in understanding from 30 what peer a route is being heard when perusing the BGP table. For example, in Table 1 all NSP A routes are assigned a local pref of 90 and all of the other routes heard from NSP A are assigned a local pref of 45. If you were to see a route tagged at a local pref of 45 in your

IBGP, that would signify a non-NSP A route announced to the customer over the customer's BGP peering with NSP A.

The net effect of this local pref hierarchy is that of the routes that we know are not NSP B or NSP A, highest local pref wins on the PNAP link. The fall-through local pref value is 5 used in the case of multiple routes heard over >1 of your connections. Multi-homed customers of the PNAP, NSP A, and NSP B would use the PNAP and, if that link was not available, the NSP A link followed by NSP B. Multi-homed customers of NSP B and NSP A would, in the example above, use NSP A followed by NSP B.

10 Whether using NSP A or NSP B in the case of a multi-homed customer of both is entirely at the customer's discretion. That behavior is easily modifiable by switching the primary and fall-through local pref sets of NSP A and NSP B.

Example 1

15 The following is an example of implementing this approach with NSP A.

NSP A peer:

```
neighbor xxx.xxx.xxx.xxx remote-as  
neighbor xxx.xxx.xxx.xxx send-community  
neighbor xxx.xxx.xxx.xxx remote-as NSP A  
20 neighbor xxx.xxx.xxx.xxx version 4  
neighbor xxx.xxx.xxx.xxx distribute-list 1 out  
neighbor xxx.xxx.xxx.xxx route-map NSP A-IN in  
neighbor xxx.xxx.xxx.xxx route-map NSP A-OUT out  
neighbor xxx.xxx.xxx.xxx filter-list 1 out  
  
25  
spr-bgw-02#
```

```
ip as-path access-list 1 permit ^$  
ip as-path access-list 2 permit .*  
30 ip as-path access-list 10 deny ^NSP A_NSP B_.*  
ip as-path access-list 10 deny ^NSP A_XXXXX_.*
```

```
route-map NSP A-OUT permit 10
! only allow customer 5 IBGP sourced routes
match as-path 1
```

5

```
route-map NSP A-IN permit 10
! let's start by denying all routes we know are NSP B and PNAP-SEA
! and attaching a medium primary local pref.
10  match as-path 10
    set local-preference 90
```

```
route-map NSP A-IN permit 20
! Any other routes attach a medium fall through local pref match as-path 2
15  set local-preference 45
```

Internap Router:

```
neighbor xxx.xxx.xxx.xxx remote-as XXXXX
neighbor xxx.xxx.xxx.xxx send-community
20 neighbor xxx.xxx.xxx.xxx version 4
neighbor xxx.xxx.xxx.xxx distribute-list 1 out
neighbor xxx.xxx.xxx.xxx route-map PNAP-IN in
neighbor xxx.xxx.xxx.xxx route-map PNAP-OUT out
neighbor xxx.xxx.xxx.xxx filter-list 1 out
25
ip community-list 1 deny 6993:NSP A      ; deny NSP A routes
ip community-list 1 deny 6993:NSP B      ; deny NSP B routes

ip as-path access-list 1 permit ^$*
30 ip as-path access-list 2 permit .*
```

```
route-map PNAP-OUT permit 10
    ! only allow customer 5 IBGP sourced routes
    ! this is already being accomplished by the distribute-list
    ! out but this routemap is where you can adjust your AS
5   ! prepends.
    match as-path 1

route-map PNAP-IN permit 10
    ! any routes that we know are not NSP B, or NSP A tag
10 highest
    ! primary local pref
    match community 1
    set local-preference 150

15 route-map PNAP-IN permit 20
    ! all else (NSP A, and NSP B routes) tag highest
    fall through
    ! local pref
    ! all else (NSP A, and NSP B routes) tag highest fall
20 through
    ! local pref
    match as-path 2
    set local-preference 75

25 NSP B Router:
    neighbor 144.228.98.5 remote-as NSP B
    neighbor 144.228.98.5 version 4
    neighbor 144.228.98.5 distribute-list 1 out
    neighbor 144.228.98.5 route-map NSP B-IN in
30    neighbor 144.228.98.5 route-map NSP B-OUT out
    neighbor 144.228.98.5 filter-list 1 out
```

```
ip as-path access-list 1 permit ^$  
ip as-path access-list 2 permit .*  
ip as-path access-list 10 deny ^NSP B_XXXXX_.*  
ip as-path access-list 10 deny ^NSP B_NSP A_.*  
5 ip as-path access-list 10 deny ^NSP B_1664_.*
```

```
route-map NSP B-OUT permit 10  
! only allow customer 5 IBGP sourced routes  
! this is already being accomplished by the distribute-list  
10 ! out but this routemap is where you can adjust your AS  
! prependings.  
match as-path 1
```

```
route-map NSP B-IN permit 10  
15 ! deny all NSP A, and PNAP routes and set a low  
! primary local pref  
match as-path 10  
set local-preference 80  
!  
20 route-map NSP B-IN permit 20  
! All else tag with a lowest fall through local pref.  
match as-path 2  
set local-preference 40
```

25 There is another configuration which also requires special consideration; namely, where a multi-homed PNAP customer with generic Diversity+ is connected to more than one PNAP.

The local-preference hierarchy of generic Diversity+ is intended to address the problem of multi-PNAP routing by creating an interlocking set of preference steps for path selection. In its default configuration, generic Diversity+ supports up to two PNAP transit
30 connections and multiple, other NSP transit connections.

Each primary level of local-preference has a corresponding secondary

value used as a backup should the primary become invalid. The complete hierarchy is shown below.

Generic Diversity + Local Preference Hierarchy (Default)

5

400 PNAP Direct Customer High (Primary Link)

350 PNAP Direct Customer Low (Secondary Link)

300 Primary PNAP Direct NSP

10 250 Secondary PNAP Direct NSP

200 Primary PNAP Non-connected

150 Secondary PNAP Non-connected

15 100 Default Local Preference Value

90 Primary PNAP Direct NSP Backup

80 Secondary PNAP Direct NSP Backup

20 70 Primary PNAP Non-connected Backup

60 Secondary PNAP Non-connected Backup

This hierarchy is applied as follows:

For customers with no more than one link to a given PNAP, routes to
25 customers of that PNAP are set to 400. When a customer has single
links to multiple PNAPs, the value is still set to 400 and the length of the AS path is left to
break the tie, meaning the direct link to the PNAP sourcing those customer routes will be used
as the AS path will be shorter.

If a customer has multiple links to the same PNAP, then routes over
30 the primary link to customers of that PNAP will be set to 400, while routes over the secondary
link to those same customer routes will be set to 350.

Routes belonging to NSPs and their customers directly connected to the primary PNAP are set to 300, while routes belonging to NSPs and their customers directly connected to the secondary PNAP are set to 250. This results in traffic being sent through the primary PNAP if the primary PNAP has a given NSP in its border fabric. If the secondary PNAP has an NSP 5 in its border fabric not common to the primary PNAP, or if an NSP common to them both fails at the primary, the traffic will be sent through the secondary for those destinations.

For destinations within NSPs which are not part of the border fabric of the primary PNAP routes are set to 200. Similar routes from the secondary PNAP are set to 150.

Should an NSP connection at the primary PNAP fail, routes to that NSP through the 10 primary PNAP will be set to 200, rather than 300. If an NSP connection at the secondary PNAP fails, routes to that NSP through the secondary PNAP will be set to 150, rather than 250.

The default value of 100 is generally not used for routes through a PNAP and is instead allocated for cases in which a customer has a connection to another NSP in addition to a 15 PNAP.

The values below 100 are used for customer NSP routes heard through the PNAP. The routes heard via the primary PNAP from the NSP to which the customer has a direct connection are set to 90. The same routes heard from the secondary PNAP are set to 80. Both of these cases assume the PNAPs have the NSP in their border fabric.

20 If the customer has a connection to an NSP not found in the border fabric of the primary PNAP, those routes heard through the primary PNAP for destinations within that NSP are set to 70. If such is the case with the secondary PNAP, those routes are set to 60.

Determining Primary and Secondary

In a simple multi-PNAP scenario, a customer is connected to more than one PNAP in 25 a given city or region and the primary and secondary PNAPs can be determined based on traffic levels within the PNAPs, provider fabric, or other concerns. However, when the multiple PNAPs are not all geographically close, a simple primary/secondary configuration may result in sub-optimal routing both in and out of the customer network.

In cases when a customer is connected to multiple, geographically diverse PNAPs the 30 preferred configuration is to have multiple primaries, one per region. In this way, PNAP NSPs will use their IGP cost for inbound traffic and the customer can similarly use their own

IGP cost for outbound traffic. Care must be taken to properly announce prefixes to control regional traffic flows. Customers with such disperse PNAP connectivity should announce both their aggregate networks as well as more specific, regional prefixes.

As an example, consider a customer with sites in both LAX and NYC with their own
5 backbone connection between them. Each site connects to one PNAP in their area. The customer has been allocated 192.168.0.0/16 and has internally allocated 192.168.0.0/17 for the LAX site and 192.168.128.0/17 for the NYC site. From the LAX PNAP they would announce both 192.168.0.0/16 and 192.168.0.0/17. From the NYC PNAP they would announce both 192.168.0.0/16 and 192.168.128.0/17. If the customer wished to avoid any
10 traffic to or from external destinations from transiting their backbone, they would instead advertise only the more specific prefixes (192.168.0.0/17 and 192.168.128.0/17) and not the aggregate (192.168.0.0/16).

This multiple primary PNAP model can be extended to an arbitrary number of regions, but within a single region, there must be a single primary.

15

Example 2

(Configuration for a Multi-PNAP Customer)

20 In the example below assume the customer is connected to two PNAPs, A and B. A is the primary, with connections to NSP C, NSP D, while B is the secondary, with connections to NSP C, NSP D, and NSP E.

PNAP Data for Customer Configuration:

25 PNAP A

Autonomous System Number: XXXXX

Border 1 Next Hop: 10.8.230.1

Internal/Customer Networks: 10.8.0.0/16

192.168.4.0/24 (AS 12005)

30 192.168.16.0/20 (AS 5507)

NSP Fabric: NSP D (AS 1239)

NSP C (AS 701)

PNAP B

| | | |
|----|-----------------------------|---|
| | Autonomous System Number: | 6993 |
| 5 | Border 2 Next Hop: | 172.18.24.33 |
| | Internal/Customer Networks: | 172.18.0.0/16 172.20.4.0/22 (AS 13461) |
| | NSP Fabric: | NSP D (AS 1239) NSP C (AS 701) |
| 10 | | NSP E (AS 3561) |

Example 3

(BGP Routes Selected at Customer)

Customer-CPE > show ip bgp

15 BGP table version is 3063602, local router ID is 10.8.230.2

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal

Origin codes: i - IGP, e - EGP, ? - incomplete

| | Network | Next Hop | Metric | LocPrf | Weight | Path |
|----|--------------------|--------------|--------|--------|--------|-------------------------|
| 20 | *>i9.2.0.0/16 | 10.8.230.1 | 0 | 300 | 0 | XXXXX XXX i |
| | * | 172.18.24.33 | 0 | 250 | 0 | 6993 XXX i |
| | *>i10.8.0.0/16 | 10.8.230.1 | 0 | 400 | 0 | XXXXX i |
| | * | 172.18.24.33 | 0 | 150 | 0 | 6993 1239 XXXXX i |
| | * 24.116.4.0/23 | 10.8.230.1 | 0 | 200 | 0 | XXXXX 1239 3561 i |
| 25 | *>i | 172.18.24.33 | 0 | 250 | 0 | 6993 3561 i |
| | *>i137.99.0.0 | 10.8.230.1 | 0 | 200 | 0 | XXXXX 1239 209 i |
| | * | 172.18.24.33 | 0 | 150 | 0 | 6993 1239 209 i |
| | * 172.18.0.0 | 10.8.230.1 | 0 | 200 | 0 | XXXXX XXX 6993 i |
| | *>i | 172.18.24.33 | 0 | 400 | 0 | 6993 i |
| 30 | * 172.20.4.0/22 | 10.8.230.1 | 0 | 200 | 0 | XXXXX XXX 6993 13461 i |
| | *>i | 172.18.24.33 | 0 | 400 | 0 | 6993 13461 i |
| | *>i192.168.4.0 | 10.8.230.1 | 0 | 400 | 0 | XXXXX 12005 i |
| | * | 172.18.24.33 | 0 | 150 | 0 | 6993 1239 XXXXX 12005 i |
| | *>i192.168.16.0/20 | 10.8.230.1 | 0 | 400 | 0 | XXXXX 5507 i |

| | | | | | |
|---|--------------|---|-----|---|------------------------|
| * | 172.18.24.33 | 0 | 150 | 0 | 6993 1239 XXXXX 5507 1 |
|---|--------------|---|-----|---|------------------------|

Detail of BGP Route Information for Specific Prefixes

Customer-CPE>sho ip bgp 10.8.0.0

5 BGP routing table entry for 10.8.0.0/16, version 1304669

Paths: (2 available, best #2)

6993 1239 XXXXX

172.18.24.33 from 172.18.24.33 (172.18.24.1)

Origin IGP, metric 0, localpref 150, valid, external

10 XXXXX

10.8.230.1 from 10.8.230.1 (10.8.230.1)

Origin IGP, metric 0, localpref 400, valid, external, best

Customer-CPE>sho ip bgp 137.99.0.0

BGP routing table entry for 137.99.0.0, version 1304669

15 Paths: (2 available, best #2)

6993 1239 209

172.18.24.33 from 172.18.24.33 (172.18.24.1)

Origin IGP, metric 0, localpref 150, valid, external

XXXXX 1239 209

20 10.8.230.1 from 10.8.230.1 (10.8.230.1)

Origin IGP, metric 0, localpref 200, valid, external, best

Although the description above contains many specificities, these should not be construed as limiting the scope of the invention but as merely providing illustrations of some 25 of the presently preferred embodiments of this invention. Thus the scope of this invention should be determined by the appended claims and their legal equivalents. Therefore, it will be appreciated that the scope of the present invention fully encompasses other embodiments which may become obvious to those skilled in the art, and that the scope of the present invention is accordingly to be limited by nothing other than the appended claims, in which 30 reference to an element in the singular is not intended to mean "one and only one" unless explicitly so stated, but rather "one or more." All structural, chemical, and functional

equivalents to the elements of the above-described preferred embodiment that are known to those of ordinary skill in the art are expressly incorporated herein by reference and are intended to be encompassed by the present claims. Moreover, it is not necessary for a device or method to address each and every problem sought to be solved by the present invention, for 5 it to be encompassed by the present claims. Furthermore, no element, component, or method step in the present disclosure is intended to be dedicated to the public regardless of whether the element, component, or method step is explicitly recited in the claims. No claim element herein is to be construed under the provisions of 35 U.S.C. 112, sixth paragraph, unless the element is expressly recited using the phrase "means for."

CLAIMS

What is claimed is:

1. A packet-switched network, comprising:
 - a Private Network Access Point (PNAP) having a customer side and a service provider side;
 - at least a first customer and a second customer connected to the customer side of the PNAP;
 - at least one service provider connected to the service provider side of the PNAP; and
 - an interconnected network system accessible to said service provider and to said first and second customers;
 - wherein traffic between said first and second customers is exchanged through the PNAP without transiting over said service provider.
2. A network as recited in claim 1, wherein the PNAP routing infrastructure within said PNAP contains a routing table listing all routes to all destinations in the network, along with a parameter indicating the degree of preference attached to each route of a set of identical routes from multiple sources.
3. A network as recited in claim 1, wherein said PNAP routing infrastructure lists a direct connection between said first and second customers within said customer side of the PNAP, and wherein said PNAP routing infrastructure sets the level of preference higher for said direct connection than for any other routes between said first and second customers.
4. A network as recited in claim 1, wherein at least one of said customers is connected to a service provider who is also connected to the service provider side of the PNAP so that said customer is multi-homed.
5. A network as recited in claim 4, wherein said multi-homed customer is provided a routing table listing all routes to all destinations in the network, along with a community attribute denoting the preferred reachability of a given route over the customer non-PNAP connection.

6. A network as recited in claim 5, wherein said multi-homed customer can utilize ASimilar data received over a BGP feed from the PNAP and send traffic to a destination served by said service provider through said customer's connection to said service provider by using said routing table to set route preferences in a router maintained by said customer.

7. A packet-switched network, comprising:

a Private Network Access Point (PNAP) having a customer side and a service provider side;

at least a first customer and a second customer connected to the customer side of the
5 PNAP;

a plurality of service providers connected to the service provider side of the PNAP; and
an interconnected network system accessible to said plurality of service providers and
to said first and second customers;

wherein the PNAP routing infrastructure within said PNAP contains a routing table
10 listing all routes to all destinations in the network, along with a parameter indicating the degree
of preference attached to each route of a set of identical routes from multiple sources;

wherein said PNAP routing infrastructure lists a direct connection between said first
and second customers within said customer side of the PNAP.

8. A network as recited in claim 7, wherein said PNAP routing infrastructure sets
the level of preference higher for said direct connection than for any other routes between said
first and second customers.

9. A network as recited in claim 7, wherein traffic between said first and second
customers is exchanged through the PNAP without transiting over said service provider.

10. A network as recited in claim 7, wherein at least one of said customers is
connected to a service provider who is also connected to the service provider side of the PNAP
so that said customer is multi-homed.

11. A network as recited in claim 10, wherein said multi-homed customer is provided with said routing table listing all routes to all destinations in the network, along with a community attribute denoting the preferred reachability of a given route over the customer non-PNAP connection.

12. A network as recited in claim 11, wherein said multi-homed customer can utilize ASimilater data received over a BGP feed from the PNAP and send traffic to a destination served by said service provider through said customer's connection to said service provider by using said routing table to set route preferences in a router maintained by said customer.

13. A method for exchanging traffic in a packet-switched network, comprising:
providing a Private Network Access Point (PNAP) having a customer side and a service provider side;

connecting at least a first customer and a second customer to the customer side of the
5 PNAP;

connecting at least one service provider to the service provider side of the PNAP;
making an interconnected network system accessible to said service provider and to said first and second customers; and

causing traffic between said first and second customers to be exchanged through the
10 PNAP without transiting over said service provider.

14. A method as recited in claim 13, wherein the PNAP routing infrastructure within said PNAP contains a routing table listing all routes to all destinations in the network, along with a parameter indicating the degree of preference attached to each route of a set of identical routes from multiple sources.

15. A method as recited in claim 13, wherein said PNAP routing infrastructure lists a direct connection between said first and second customers within said customer side of the PNAP.

16. A method as recited in claim 15, wherein said PNAP routing infrastructure sets the level of preference higher for said direct connection than for any other routes between said first and second customers.

17. A method as recited in claim 13, wherein at least one of said customers is connected to a service provider who is also connected to the service provider side of the PNAP so that said customer is multi-homed.

18. A method as recited in claim 17, further comprising providing said multi-homed customer with a routing table listing all routes to all destinations in the network, along with a community attribute denoting the preferred reachability of a given route over the customer non-PNAP connection.

19. A method as recited in claim 18, further comprising allowing said multi-homed customer to utilize ASimilater data received over the BGP feed from the PNAP and send traffic to a destination served by said service provider through said customer's connection to said service provider by using said routing table to set route preferences in a router maintained by
5 said customer.

20. A method for exchanging traffic in a packet-switched network, comprising:
providing a Private Network Access Point (PNAP) having a customer side and a service provider side;

connecting at least a first customer and a second customer to the customer side of the
5 PNAP;

connecting at least one service provider to the service provider side of the PNAP; and
making an interconnected network system accessible to said service provider and to said first and second customers;

wherein the PNAP routing infrastructure within said PNAP contains a routing table
10 listing all routes to all destinations in the network, along with a parameter indicating the degree of preference attached to each route of a set of identical routes from multiple sources;

wherein said PNAP routing infrastructure lists a direct connection between said first

and second customers within said customer side of the PNAP.

21. A method as recited in claim 20, further comprising causing traffic between said first and second customers to be exchanged through the PNAP without transiting over said service provider.

22. A method as recited in claim 20, wherein said PNAP routing infrastructure sets the level of preference higher for said direct connection than for any other routes between said first and second customers.

23. A method as recited in claim 20, wherein at least one of said customers is connected to a service provider who is also connected to the service provider side of the PNAP so that said customer is multi-homed.

24. A method as recited in claim 23, further comprising providing said multi-homed customer with said routing table listing all routes to all destinations in the network, along with a community attribute denoting the preferred reachability of a given route over the customer non-PNAP connection.

25. A method as recited in claim 24, further comprising allowing said multi-homed customer to utilize ASimilater data received over a BGP feed from the PNAP and send traffic to a destination served by said service provider through said customer's connection to said service provider by using said routing table to set route preferences in a router maintained by
5 said customer.

26. A packet-switched network, comprising:
a Private Network Access Point (PNAP) having a customer side and a service provider side;
at least one customer connected to the customer side of the PNAP;
5 at least one service provider connected to the service provider side of the PNAP; and
an interconnected network system accessible to said service provider and to said

customer;

wherein said customer is connected to a service provider who is also connected to the service provider side of the PNAP so that said customer is multi-homed;

10 wherein said multi-homed customer is provided a routing table listing all routes to all destinations in the network, along with a community attribute denoting the preferred reachability of a given route over the customer non-PNAP connection; and

wherein said multi-homed customer can utilize the ASimilater data received over the BGP feed from the PNAP and send traffic to a destination served by said service provider
15 through said customer's connection to said service provider by using said routing table to set route preferences in a router maintained by said customer.

27. A method for exchanging traffic in a packet-switched network, comprising:
providing a Private Network Access Point (PNAP) having a customer side and a service provider side;

5 connecting at least one customer to the customer side of the PNAP;
connecting at least one service provider to the service provider side of the PNAP;
making an interconnected network system accessible to said service provider and to said customer;

wherein said customer is connected to a service provider who is also connected to the service provider side of the PNAP so that said customer is multi-homed;

10 providing said multi-homed customer a routing table listing all routes to all destinations in the network, along with a community attribute denoting the preferred reachability of a given route over the customer non-PNAP connection; and

allowing said multi-homed customer to utilize the ASimilater data received over the BGP feed from the PNAP and send traffic to a destination served by said service provider
15 through said customer's connection to said service provider by using said routing table to set route preferences in a router maintained by said customer.

28. A PNAP system, comprising:

- a) a plurality of packet switched routing control devices;
- b) a private network access point (PNAP) customer connected to a first side of one

- of said packet switched routing control devices; and
- 5 c) a plurality of network service providers connected to a second side of said one of said packet switched routing control devices, said one of said packet switched routing control devices causing packets originated by said PNAP customer to enter into said one of said packet switched routing control devices from said first side and be routed directly to a selected one of said plurality of network
- 10 service providers, said selected one of said network service providers containing a destination network which the PNAP customer desires to communicate with said packet switched routing control devices causing a response packet originated from the destination network to be forwarded from the destination network back to the originator PNAP customer over the exact
- 15 same path, thus providing symmetric routing.
29. A PNAP system as recited in claim 28, wherein said plurality of network service providers have local preferences set for causing said network service providers to select direct routing to said one of said packet switched routing control devices for causing response packets to be routed over the exact same return path as a packet received over a forward path.
- 5 30. A method for communicating employing a private network access point (PNAP) system, comprising:
- a) causing packets originated by a private network access point (PNAP) customer to enter into one of a plurality of packet switched routing control devices;
- 5 b) causing entering packets to be routed directly to a provider which contains a desired destination network; and
- c) causing a response packet from said destination network to be routed back to the PNAP customer over the exact same path.
31. A method as recited in claim 30, further comprising:
- a) setting local preferences of the provider to cause the provider to select direct routing back to the one of a plurality of packet switched routing control devices.

32. A plurality of packet switched routing control devices, comprising:
- a) means for causing packets originated by a private network access point (PNAP) customer to enter into at least one of said packet switched routing devices;
 - b) means for causing entering packets to be routed directly to a provider which contains a desired destination network; and
 - c) means for causing a response packet from said destination network to be routed directly back to the PNAP customer over the exact same path.
- 5
33. A plurality of packet switched routing devices as recited in claim 32, further comprising:
- d) means for setting local preferences of the provider to cause the provider to select direct routing back to the at least one of said packet switched routing devices.
- 5
34. A method for routing information packets symmetrically on selected forward and return paths in an overall network involving a plurality of traffic carrying networks, said overall network including a plurality of public network access point (NAPs) and a plurality of private network access points (PNAPs), said plurality of traffic carrying networks including non-PNAP providers, and PNAP providers, each of said traffic carrying networks having Autonomous System (AS) numbers associated therewith, said method comprising:
- a) creating a list of all PNAP provider AS numbers;
 - b) creating a list of AS numbers which peer at the plurality of public NAPs but which are not any of the AS numbers associated with said PNAP providers;
 - c) for each provider, taking the union of all provider AS numbers and AS numbers associated with the plurality of public NAPs and subtracting out the AS numbers associated with the current provider;
 - d) denying that list of AS numbers for the current provider;
 - e) setting said local preferences of the PNAP providers for causing the PNAP providers to select direct routing to the plurality of PNAPs; and
- 10
- 15

- f) applying configuration files to routers associated with each PNAP provider; and causing the routers to apply the PNAP provider configuration files to the AS numbers received from each provider.
35. An apparatus for routing information packets symmetrically on selected forward and return paths in an overall network involving a plurality of traffic carrying networks, said overall network including a plurality of public network access point (NAPs) and a plurality of private network access points (PNAPs), said plurality of traffic carrying networks including non-PNAP providers and PNAP providers, each of said traffic carrying networks having Autonomous System (AS) numbers associated therewith, said apparatus comprising:
- a) means for creating a list of all PNAP provider AS numbers;
 - b) means for creating a list of AS numbers which peer at the plurality of public NAPs but which are not any of the AS numbers associated with said PNAP providers;
 - c) means for taking the union of all provider AS numbers and AS numbers associated with the plurality of public NAPs and subtracting out the AS numbers associated with the current provider;
 - d) means for denying that list of AS numbers for the current provider;
 - e) means for setting local preferences of the PNAP providers for causing the PNAP providers to select direct routing to the plurality of PNAPs;
 - f) means for applying configuration files to routers associated with each PNAP provider; and
 - g) means for causing the routers to apply the PNAP provider configuration files to the AS numbers received from each provider.
- 20
36. A method for routing information packets symmetrically on selected forward and return paths in an overall network involving a plurality of traffic carrying networks, said overall network including a plurality of public network access points (NAPs) and a plurality of private network access points (PNAPs), said plurality of traffic carrying networks being divided into two groups, a first group consisting of non-PNAP
- 5

providers, and a second group consisting of PNAP providers, each of said traffic carrying networks having one or more Autonomous System (AS) numbers associated therewith, said method comprising:

- a) creating a list of all PNAP provider AS numbers;
- 10 b) creating a list of AS numbers which peer at the plurality of public NAPs but which are not any of the AS numbers associated with said PNAP providers;
- c) for each provider, taking the union of all provider AS numbers and AS numbers associated with the plurality of public NAPs and subtracting out the AS numbers associated with the current provider;
- 15 d) denying that list of AS numbers for the current provider;
- e) tagging all other AS numbers which do not match the "deny" list with a primary preference value associated with the provider;
- f) tagging all AS numbers which do match the "deny" list with a secondary preference value associated with the provider;
- 20 g) using the primary local preference values to cause the plurality of PNAPs to route to the PNAP provider's destinations over that provider's network;
- h) using the secondary local preference values to cause the plurality of PNAPs to route to destinations not connected to a PNAP provider over a preselected PNAP provider;
- 25 i) making changes to AS path lengths of routes advertised by the plurality of PNAPs to each PNAP provider to cause providers not directly connected to the PNAP to use the same preselected PNAP provider as the plurality of PNAPs uses to send to said providers; and
- j) setting the PNAP provider local preferences for causing the PNAP provider to select direct routing to the plurality of PNAPs;
- 30 k) applying configuration files to routers associated with each PNAP provider; and causing the routers to apply the PNAP provider configuration files to the routes received from each provider.

37. A private network access point (PNAP) system, comprising:

- a) a plurality of packet switched routing control devices;

- b) a private network access point (PNAP) customer connected to a first side of one of said plurality of packet switched routing control devices;
- c) a plurality of PNAP network service providers directly connected to a second side of said plurality of packet switched routing control devices; and
- 5 d) at least one non-PNAP network service provider not directly connected to said plurality of said packet switched routing control devices, said packet switched routing control devices causing packets originated by said private PNAP customer to enter into said packet switched routing control devices from said first side, said packet switched routing control devices routing packets addressed to a destination within said non-PNAP network service provider's networks to a preselected one of said plurality of PNAP network service providers, said preselected one of said plurality of PNAP network service providers containing a route to the non-PNAP destination network which the private PNAP customer desires to communicate with, and the packet switched routing control devices causing a response packet originated from the destination to be forwarded from the destination back to the private PNAP customer over the exact same path, thus providing symmetric routing.
- 10
- 15
38. A PNAP system as recited in claim 37, in which said plurality of PNAP network service providers have local preferences set for causing said PNAP network service providers to select direct routing to said packet switched routing control devices for causing response packets to be routed over the exact same return path as a packet received over a forward path.
- 5
39. A method for communicating employing a private network access point (PNAP) system, comprising:
- a) causing packets originated by a private network access point (PNAP) customer to enter into one of a plurality of packet switched routing control devices;
- 5 b) causing entering packets addressed to a destination within a non-PNAP network service provider's networks to be routed to a preselected one of a plurality of PNAP network service providers containing a route to the non-PNAP

- destination; and
- 10 c) causing a response packet originated from the destination to be forwarded from the destination back to the private PNAP customer over the exact same path.
40. A method as recited in claim 39, further comprising:
- d) setting local preferences of the PNAP network service providers to cause the PNAP network service providers to select direct routing back to said one of said plurality of said packet switched routing control devices.
41. A plurality of packet switched routing control devices, comprising:
- a) means for causing packets originated by a private network access point (PNAP) customer to enter into one of said plurality of said packet switched routing control devices;
- 5 b) means for causing entering packets addressed to a destination within a non-PNAP network service provider's networks to be routed to a preselected one of a plurality of PNAP network service providers containing a route to the non-PNAP destination; and
- c) means for causing a response packet originated from the destination to be forwarded from the destination back to the private PNAP customer over the exact same path.
- 10 42. The plurality of packet switched routing control devices as recited in claim 41, further comprising:
- a) means for setting local preferences of the plurality of PNAP network service providers to cause the PNAP network providers to select direct routing back to said one of said plurality of packet switched routing devices.
- 15 43. An apparatus, comprising a plurality of private network access point (PNAP), a provider having a connection to said plurality of PNAPs, said provider having networks, said networks having destinations, said apparatus configured to route packets symmetrically between destinations and said plurality of PNAPs, within the provider's networks, over that provider's connection to the plurality of PNAPs.

44. A method for routing symmetrically between destinations within a private network access point (PNAP) provider's networks, over that PNAP provider's connection to a plurality of PNAPs, said method comprising: applying configuration files to routers associated with PNAP provider's networks; and causing the routers to apply the configuration files to routes received from each provider.
5
45. An apparatus, comprising a plurality of private network access points (PNAPs), a PNAP provider having a connection to said plurality of PNAPs, a non-PNAP provider not having a direct connection to said plurality of PNAPs, said providers having networks, said non-PNAP provider's networks having destinations, said apparatus to route symmetrically to/from destinations within non-PNAP provider's networks over a selected PNAP provider's connection to the plurality of PNAPs.
10
46. A method for routing symmetrically between destinations within a non-private network access point (non-PNAP) provider's networks over a selected PNAP provider's connection to a plurality of PNAPs, said method comprising: applying configuration files to routers associated with PNAP provider's networks; and causing the routers to apply the configuration files to routes received from each provider.
5
47. An apparatus, comprising a plurality of private network access points (PNAPs), a first and second provider having connections to said plurality of PNAPs, said first and second provider having networks, said first provider's networks having a destination, said apparatus configured for routing symmetrically to/from destinations within the first PNAP provider's networks over the second PNAP provider's connection to the plurality of PNAPs when the first PNAP provider's connection to the plurality of PNAPs is unavailable.
5
48. A method, comprising routing symmetrically between destinations within a private network access point (PNAP) provider's networks over the next best PNAP provider's

connection to the plurality of PNAPs when the destination's PNAP provider connection to the plurality of PNAPs is unavailable.

49. A network topology, comprising:
 - two or more network service providers;
 - two or more private network access points (PNAPs); and
 - one or more customers,

5 whereby the topology is used to exchange packets between the customer and the providers in a symmetric fashion.
50. A network topology, comprising:
 - two or more network service providers;
 - two or more private network access points (PNAPs); and
 - one or more customers,

5 whereby the topology is used exclusively to provide for the symmetric exchange of packets between a customer and a network service provider.
51. A network topology as recited in claim 50, wherein service is provided for the symmetric exchange of packets between a private customer and a network service provider.
52. A packet switched routing control apparatus for the symmetric routing of packets, comprising:
 - a) means for determining symmetric routing policy at a plurality of private network access points (PNAPs);
 - b) means for creating and maintaining router configurations in accordance with said routing policy;
 - c) means for causing a PNAP provider to set a preference to prefer access to the plurality of PNAPs via a local connection; and
 - d) means for causing other providers to prefer access to the plurality of PNAPs in accordance with said symmetric routing policy.

10

53. A packet switched routing control method for the symmetric routing of packets, comprising:
- a) determining symmetric routing policy at a plurality of private network access points (PNAPs);
 - 5 b) creating and maintaining router configurations in accordance with said routing policy;
 - c) causing PNAP providers to set a preference to prefer access to the plurality of PNAPs via a local connection; and
 - d) causing other providers to prefer access to the plurality of PNAPs in accordance with said symmetric routing policy.
- 10
54. A packet switched routing control apparatus for the symmetric routing of packets, comprising:
- a) means for determining a LOCAL_PREF ordering of private network access point (PNAP) providers to be used when a destination network is not connected to a PNAP provider or the destination network is not currently reachable through its PNAP provider or the destination network is connected to multiple PNAP providers;
 - 5 b) means for determining which Autonomous Systems (ASes) are directly associated with each PNAP provider and storage of those Autonomous Systems in a Provider AS Database;
 - c) means for determining which Autonomous Systems are peering at a plurality of public NAPs and storage of those Autonomous Systems in an Exception AS Database;
 - 10 d) means for determining the Autonomous System (AS) numbers of other providers provided to the apparatus by the user and storage of those Autonomous Systems in the Provider AS Database;
 - e) means for verifying traffic flows from a plurality of PNAPs to said PNAP providers and updating databases;
 - 15 f) means for verifying traffic flows from said PNAP providers to the plurality of

- 20 PNAPs and creating appropriate notifications;
- g) means for creating base router configuration files;
- h) means for adding LOCAL_PREF configuration commands to said base router configuration files such that packets from the plurality of PNAPs to a destination within said PNAP provider's network, traverse that provider's connection to the plurality of PNAPs;
- 25 i) means for adding LOCAL_PREF configuration commands to said base router configuration files such that packets from the plurality of PNAPs to a destination not within said PNAP provider's network, traverse the PNAP provider's network with the most preferred preference;
- 30 j) a computing means for determining the appropriate AS_PATH additions to be added to the routes sent from the plurality of PNAPs to each PNAP provider and storage of those additions in the Provider AS_PATH Prepend Database;
- k) means for adding AS_PATH Prepend configuration commands to said base router configuration files using said AS_PATH Prepend Database;
- 35 l) means for causing other providers to prefer access back to the plurality of PNAPs over a preferred PNAP provider's network;
- m) means for applying said combined router configuration files to routers connected to each PNAP provider;
- n) means for causing each router to obtain full routes from each provider; and
- 40 o) means for causing each router to apply said router configurations to said full routes.
55. A packet switched routing control method for the symmetric routing of packets, comprising:
- 5 a) determining a LOCAL_PREF ordering of a plurality of private network access point (PNAP) providers to be used when a destination network is not connected to a PNAP provider or the destination network is not currently reachable through its PNAP provider or the destination network is connected to multiple PNAP providers;
- b) determining which Autonomous Systems (ASes) are directly associated with

- each PNAP provider and storage of those Autonomous Systems in the Provider AS Database;
- 10 c) determining which Autonomous Systems are peering at a plurality of public NAPs and storage of those Autonomous Systems in the Exception AS Database;
- d) determining the Autonomous System (AS) numbers of other providers provided to the system by the user and storage of those Autonomous Systems in the Provider AS Database comprising the steps of;
- 15 e) verifying traffic flows from the plurality of PNAPs to said PNAP providers and updating databases;
- f) verifying traffic flows from said PNAP providers to a plurality of PNAPs creating appropriate notifications;
- 20 g) creating base router configuration files;
- h) adding a plurality of LOCAL_PREF configuration commands to said base router configuration files such that said packets from the plurality of PNAPs to a destination within a PNAP provider's network, traverse that said provider's connection to the plurality of PNAPs;
- 25 i) adding said LOCAL_PREF configuration commands to said base router configuration files such that said packets from the plurality of PNAPs to a destination not within a PNAP provider's network, traverse the PNAP provider's network with the most preferred preference;
- j) determining the appropriate AS_PATH additions to be added to the routes sent from the plurality of PNAPs to each PNAP provider and storage of those additions in the Provider AS_PATH Prepend Database;
- 30 k) adding a plurality of AS_PATH Prepend configuration commands to said base router configuration files using said AS_PATH Prepend Database;
- l) causing other providers to prefer access back to the plurality of PNAPs over a preferred PNAP provider's network;
- 35 m) applying combined router configuration files to routers connected to each PNAP provider;
- n) causing each router to obtain full routes from each provider; and
- o) causing each router to apply said router configurations to said full routes.

- 40 56. An apparatus, including a plurality of private network access points (PNAPs), for determining a LOCAL_PREF ordering of PNAP providers to be used when a destination network is not connected to a PNAP provider or the destination network is not currently reachable through its PNAP provider or the destination network is connected to multiple PNAP providers, comprising:
- 45 a) means for creating a first set (primary) of descending order LOCAL_PREF values, one for each PNAP provider, such that the highest LOCAL_PREF value is the most preferred PNAP provider; and
- b) means for creating a second set (secondary) of LOCAL_PREF values, one for each PNAP provider, in the same descending order, such that the highest LOCAL_PREF value in the second set is a value lower than the lowest LOCAL_PREF value of the first set.
- 50
- 55 57. A method for determining a LOCAL_PREF ordering of PNAP providers to be used when a destination network is not connected to a PNAP provider or the destination network is not currently reachable through its PNAP provider or the destination network is connected to multiple PNAP providers, comprising the steps of:
- a) creating a first set (primary) of descending order LOCAL_PREF values, one for each PNAP provider, such that the highest LOCAL_PREF value is the most preferred PNAP provider; and
- 60 b) creating a second set (secondary) of LOCAL_PREF values, one for each PNAP provider, in the same descending order, such that the highest LOCAL_PREF value in the second set is a value lower than the lowest LOCAL_PREF value of the first set.
58. An apparatus for determining whether an Autonomous System of a plurality of Autonomous Systems is directly associated with each PNAP provider, comprising:
- a) a means for initially priming a Provider AS Database as <AS, Provider> pairs using the directly connected AS number of each Provider;
- 5 b) a means for loading the existing Provider AS Database as <AS, Provider> pairs;

- c) a means for searching through each AS and using the whois mechanism, determining if that AS continues to correspond to that AS's Provider;
 - d) a means for removing any ASes from the Provider AS Database that are determined to no longer correspond to their Provider;
 - e) a means for downloading and processing a full routing table dump from each Provider as a list of AS_PATH attributes for each route;
 - f) a computing means for taking said AS_PATH lists and using them to determine additional AS numbers associated with each PNAP provider and storage of the <AS, Provider> pairs in the Provider AS Database;
 - g) a computing means for taking said AS_PATH lists and using them to determine AS numbers associated with other Providers as provided by the user and storage of the <AS, Provider> pairs in the Provider AS Database.
59. A method for determining whether an Autonomous System of a plurality of Autonomous Systems is directly associated with each PNAP provider, said method comprising the steps of:
- a) initially priming a Provider AS Database as <AS, Provider> pairs using the directly connected AS number of each Provider;
 - b) loading the existing Provider AS Database as <AS, Provider> pairs;
 - c) searching through each AS and using the whois mechanism, determining if that AS continues to correspond to that AS's Provider;
 - d) removing any ASes from the Provider AS Database that are determined to no longer correspond to their Provider;
 - e) downloading and processing a full routing table dump from each Provider as a list of AS_PATH attributes for each route;
 - f) computing means for taking said AS_PATH lists and using them to determine additional AS numbers associated with each PNAP provider and storage of the <AS, Provider> pairs in the Provider AS Database; and
 - g) taking said AS_PATH lists and using them to determine AS numbers associated with other Providers as provided by the user and storage of the <AS, Provider> pairs in the Provider AS Database.

60. An apparatus, for taking a plurality of AS_PATH lists and using them to determine a plurality of additional AS numbers associated with each PNAP provider and a plurality of those AS numbers associated with Providers provided by the user, comprising:
- a) for each AS within the AS_PATH, a means for moving from left to right in the path;
 - b) a means for determining if this is the first AS in the path and if so lookup the AS in the Provider AS Database to determine the Provider advertising this AS_PATH and saving that Provider for future additions to the Provider AS Database;
 - c) a means for using whois to lookup the AS information for any subsequent AS numbers;
 - d) a means for determining if the AS information corresponds to the same Provider advertising this AS_PATH and if so add this <AS, Provider> pair to the AS Database;
 - e) a means for stopping the method when you reach an AS with AS information which does not correspond to the same Provider advertising the AS_PATH;
 - f) a means for determining if the said AS information corresponds to a user provided Provider and if so add this <AS, user-provided Provider> pair to the AS Database and end the method; and
 - g) a means for incrementing a counter for those AS numbers whose AS information does not correspond to the same Provider advertising this AS_PATH or Provider provided by the user.
61. A method for taking a plurality of AS_PATH lists and using them to determine a plurality of additional AS numbers associated with each PNAP provider and a plurality of those AS numbers associated with Providers provided by the user, comprising the steps of:
- a) for each AS within the AS_PATH, moving from left to right in the path;
 - b) determining if this is the first AS in the path and if so lookup the AS in the Provider AS Database to determine the Provider advertising this AS_PATH and

- saving that Provider for future additions to the Provider AS Database;
- c) using whois to lookup the AS information for any subsequent AS numbers;
- d) determining if the said AS information corresponds to the same Provider advertising this AS_PATH and if so add this <AS, Provider> pair to the AS Database;
- e) stopping the method when you reach an AS with AS information which does not correspond to the same Provider advertising the AS_PATH;
- f) determining if the said AS information corresponds to a user provided Provider and if so add this <AS, user-provided Provider> pair to the AS Database and end the method; and
- g) incrementing a counter for those AS numbers whose AS information does not correspond to the same Provider advertising this AS_PATH or Provider provided by the user.
- 20
62. An apparatus for taking a plurality of AS_PATH lists and using them to determine a plurality of additional AS numbers associated with each PNAP provider and a plurality of those AS numbers associated with Providers provided by the user, comprising:
- a) a means for determining the first AS in the path looking up the AS in the Provider AS Database to determine the Provider advertising this AS_PATH and saving that Provider for future additions to the Provider AS Database;
- b) for each AS within the AS_PATH, a means for moving from right to left in the path;
- c) a means for using whois to lookup the AS information for any subsequent AS numbers;
- d) a means for determining if the said AS information corresponds to the same Provider advertising this AS_PATH and if so add this <AS, Provider> pair to the AS Database;
- e) a means for stopping the method when you reach the end of the AS_PATH;
- f) a means for determining if the AS information corresponds to a user provided Provider and if so add this <AS, user-provided Provider> pair to the AS Database and end the method; and

- g) a means for incrementing a counter for those AS numbers whose AS information does not correspond to the same Provider advertising this AS_PATH or Provider provided by the user.
- 20
63. A method for taking a plurality of AS_PATH lists and using them to determine a plurality of additional AS numbers associated with each PNAP provider and a plurality of those AS numbers associated with Providers provided by the user, comprising the steps of:
- 5 a) determining the first AS in the path looking up the AS in the Provider AS Database to determine the Provider advertising this AS_PATH and saving that Provider for future additions to the Provider AS Database;
- b) for each AS within the AS_PATH, moving from right to left in the path;
- c) using who is to lookup the AS information for any subsequent AS numbers;
- 10 d) determining if the said AS information corresponds to the same Provider advertising this AS_PATH and if so add this <AS, Provider> pair to the AS Database;
- e) stopping the method when you reach the end of the AS_PATH;
- f) determining if the said AS information corresponds to a user provided Provider and if so add this <AS, user-provided Provider> pair to the AS Database and end the method; and
- 15 g) incrementing a counter for those AS numbers whose AS information does not correspond to the same Provider advertising this AS_PATH or Provider provided by the user.
- 20
64. An apparatus for determining whether an Autonomous System of a plurality of Autonomous Systems is peering at a plurality of Public NAPs, comprising:
- a) a means for determining those AS numbers with a counter greater than 3; and
- b) a means for adding said AS numbers to an Exception AS Database.
65. A method for determining whether an Autonomous System of a plurality of Autonomous Systems is peering at a plurality of Public NAPs, comprising the steps of:

- a) determining those AS numbers with a counter greater than 3; and
 - b) adding said AS numbers to an Exception AS Database.
66. An apparatus for verifying traffic flows from a plurality of PNAPs to a plurality of PNAP providers and updating databases if verification fails, comprising:
- a) a means for locating a traceroute server within each of the PNAP provider networks;
 - 5 b) a means for running prtraceroute to each PNAP provider traceroute server resulting in an ordered list of Autonomous System numbers for each PNAP provider;
 - c) a means for determining for each list if more than one PNAP provider is represented in the list of AS numbers from prtraceroute and if the Provider connection to the plurality of PNAPs is currently active then;
 - 10 d) a means for searching through the AS list from the destination AS back to the PNAP AS (right to left) to locate the first other PNAP provider in the list then;
 - e) a means for checking the Exception AS Database for the AS just to the left (previous) of the first other PNAP provider in the list and removing it if it exists because it was incorrectly added; and
 - 15 f) a means for sending e-mail to interested parties if the AS just to the left (previous) of the first other PNAP provider in the list, exists but is not in the Exception AS Database.
67. A method for verifying traffic flows from the plurality of PNAPs to a plurality of PNAP providers and updating databases if verification fails, comprising the steps of:
- a) locating a traceroute server within each of the PNAP provider networks;
 - b) running prtraceroute to each PNAP provider traceroute server resulting in an ordered list of Autonomous System numbers for each PNAP provider;
 - 5 c) determining for each list if more than one PNAP provider is represented in the list of AS numbers from prtraceroute and if the Provider connection to the PNAP is currently active then;
 - d) searching through the AS list from the destination AS back to the PNAP AS

- 10 (right to left) to locate the first other PNAP provider in the list then;
- e) checking the Exception AS Database for the AS just to the left (previous) of the first other PNAP provider in the list and removing it if it exists because it was incorrectly added; and
- f) sending e-mail to interested parties if the AS just to the left (previous) of the first other PNAP provider in the list, exists but is not in the Exception AS Database.
- 15
68. An apparatus for verifying traffic flows from a plurality of PNAP providers to the plurality of PNAPs and creating appropriate notifications if verification fails, comprising:
- a) a means for locating a traceroute server within each of the PNAP provider networks;
- b) a means for running traceroute from each server back to the plurality of PNAPs and running that resulting output through a modified prtraceroute resulting in an ordered list of Autonomous System numbers for each PNAP provider;
- c) a means for determining for each list if more than one PNAP provider is represented in the list of AS numbers and if the Provider connection to the plurality of PNAPs is currently active then; and
- d) a means for sending notification of a routing anomaly.
- 10
69. A method for verifying traffic flows from a plurality of PNAP providers to a plurality of PNAPs and creating appropriate notifications if verification fails, comprising the steps of:
- a) locating a traceroute server within each of a plurality of PNAP provider networks;
- b) running traceroute from each server back to the plurality of PNAPs and running that resulting output through a modified prtraceroute resulting in an ordered list of Autonomous System numbers for each PNAP provider;
- c) determining for each list if more than one PNAP provider is represented in the list of AS numbers and if the Provider connection to the plurality of PNAPs is

- currently active then; and
- d) sending notification of a routing anomaly.
70. An apparatus for creating base router configuration files for a plurality of PNAP providers, comprising:
- a) a means for creating a command describing the IP address and Autonomous System of PNAP provider;
- 5 b) a means for creating a command describing the version (4) of BGP;
- c) a means for creating a command describing the base mechanism to use in attaching LOCAL_PREF values to routes received from the PNAP provider; and
- d) a means for creating a command describing the base mechanism to use in attaching additional Autonomous System numbers to PNAP routes sent to the PNAP provider.
- 10
71. A method for creating base router configuration files for a plurality of PNAP providers, comprising the steps of:
- a) creating a command describing the IP address and Autonomous System of PNAP provider;
- 5 b) creating a command describing the version (4) of BGP;
- c) creating a command describing the base mechanism to use in attaching LOCAL_PREF values to routes received from the PNAP provider; and
- d) creating a command describing the base mechanism to use in attaching additional Autonomous System numbers to PNAP routes sent to the PNAP provider.
- 10
72. An apparatus for adding LOCAL_PREF configuration commands to base router configuration files such that packets from a plurality of PNAPs to a destination within PNAP provider's network, traverse that Provider's connection to a plurality of PNAPs, comprising:
- 5 a) a means for determining that for each PNAP provider a list is created of all

- other PNAP providers;
- b) a means for searching a provider AS Database for AS numbers of all other PNAP providers and combining that with all AS numbers contained in an Exception AS Database;
- 10 c) a means for creating a filter denying receipt of routes containing each said AS number and permitting receipt of all other routes;
- d) a means for applying a LOCAL_PREF value to permitted routes taken from the primary LOCAL_PREF values list for each Provider;
- e) a means for creating a filter permitting all routes previously denied; and
- 15 f) a means for applying a LOCAL_PREF value to said permitted routes taken from the secondary LOCAL_PREF values list for each provider.
73. A method for adding LOCAL_PREF configuration commands to base router configuration files such that packets from a plurality of PNAPs to a destination within a PNAP provider's network, traverse that Provider's connection to the plurality of PNAPs, comprising the steps of:
- 5 a) determining that for each PNAP provider a list is created of all other PNAP providers;
- b) searching a Provider AS Database for AS numbers of all other PNAP providers and combining that with all AS numbers contained in an Exception AS Database;
- 10 c) creating a filter denying receipt of routes containing each said AS number and permitting receipt of all other routes;
- d) applying the LOCAL_PREF value to permitted routes taken from the primary LOCAL_PREF values list for each Provider;
- e) creating a filter permitting all routes previously denied; and
- 15 f) applying the LOCAL_PREF value to permitted routes taken from the secondary LOCAL_PREF values list for each Provider.
74. An apparatus for adding LOCAL_PREF configuration commands to base router configuration files such that packets from a plurality of PNAPs to a destination not

- 20 within a PNAP provider's network, traverse the PNAP provider's network with the most preferred preference, comprising:
- a) a means for determining that for each PNAP provider a list is created of all other PNAP providers;
 - b) a means for searching a Provider AS Database for AS numbers of all other PNAP providers and combining that with all AS numbers contained in an Exception AS Database;
 - c) a means for creating a filter denying receipt of routes containing each said AS number and permitting receipt of all other routes;
 - d) a means for applying the LOCAL_PREF value to said permitted routes taken from the primary LOCAL_PREF values list for each Provider;
 - e) a means for creating a filter permitting all routes previously denied; and
 - f) a means for applying the LOCAL_PREF value to said permitted routes taken from the secondary LOCAL_PREF values list for each Provider.
75. A method for adding LOCAL_PREF configuration commands to base router configuration files such that packets from a plurality of PNAPs to a destination not within a PNAP provider's network, traverse the PNAP provider's network with the most preferred preference, comprising the steps of:
- a) determining that for each PNAP provider a list is created of all other PNAP providers;
 - b) searching a Provider AS Database for AS numbers of all other PNAP providers and combining that with all AS numbers contained in an Exception AS Database;
 - c) creating a filter denying receipt of routes containing each said AS number and permitting receipt of all other routes;
 - d) applying a LOCAL_PREF value to said permitted routes taken from the primary LOCAL_PREF values list for each Provider;
 - e) creating a filter permitting all routes previously denied; and
 - f) applying the LOCAL_PREF value to said permitted routes taken from the secondary LOCAL_PREF values list for each Provider.

76. A method for determining appropriate AS_PATH additions to be added to routes advertised from a plurality of PNAPs to PNAP provider's and storage of those additions in a Provider AS_PATH Prepend Database, the method, comprising the steps of:
- a) retrieving, from a plurality of public NAPs, a sample PNAP route and the route's associated AS_PATH from each PNAP provider and storing the length of said AS_PATH in a Provider AS_PATH Prepend Database;
 - b) retrieving, for each PNAP provider, the AS_PATH length (PAPL0) from the AS_PATH Prepend Database;
 - c) retrieving, for each other PNAP provider whose AS_PATH length value has never been associated with PAPL0, the AS_PATH length (PAPL1) from the AS_PATH Prepend Database;
 - d) comparing PAPL0 with PAPL1, and if PAPL0 is greater than or equal to PAPL1, then increase PAPL1 by a value of one (1) and re-storing the value back in the AS_PATH Prepend Database;
 - e) continuing to compare through all other PNAP providers (PAPL1); and
 - f) continuing to compare through all PNAP providers (PAPL0).
77. A commands to base method of adding AS_PATH Prepend configuration router configuration files, comprising the steps of:
- a) retrieving, for each PNAP provider, the AS_PATH length from an AS_PATH Prepend Database; and
 - b) creating a router configuration command which will increase the AS_PATH length of routes advertised to each PNAP provider by the value retrieved from the AS_PATH Prepend Database.
78. A method for determining appropriate AS_PATH additions to be added to routes advertised from a plurality of PNAPs to each PNAP provider and storage of those additions in a Provider AS_PATH Prepend Database, the method comprising the steps of:

- 5 a) retrieving, from a plurality of public NAPs, a sample PNAP route and the route's associated AS_PATH from each PNAP provider and storing the length of said AS_PATH in a Provider AS_PATH Prepend Database;
- b) retrieving, for each PNAP provider, its corresponding AS_PATH length (PAPL0) from an AS_PATH Prepend Database;
- 10 c) retrieving, for each other PNAP provider whose AS_PATH length value has never been associated with PAPL0, the AS_PATH length (PAPL1) from the AS_PATH Prepend Database;
- d) comparing PAPL0 with PAPL1, and if PAPL0 is greater than or equal to PAPL1, then increase PAPL1 by a value of one (1) and re-storing the value back in the AS_PATH Prepend Database;
- 15 e) continuing to compare through all other PNAP providers (PAPL1); and
- f) continuing to compare through all PNAP providers (PAPL0).
79. An apparatus for adding AS_PATH Prepend configuration commands to base router configuration files comprising:
- a) a means for retrieving for each PNAP provider the AS_PATH length from an AS_PATH Prepend Database; and
- 5 b) a means for creating a router configuration command which will increase the AS_PATH length of routes advertised to each PNAP provider by the value retrieved from the AS_PATH Prepend Database.
80. A method for adding AS_PATH Prepend configuration commands to base router configuration files, comprising the steps of:
- a) retrieving, for each PNAP provider, the AS_PATH length from an AS_PATH Prepend Database; and
- 5 b) creating a router configuration command which will increase the AS_PATH length of routes advertised to each PNAP provider by the value retrieved from the AS_PATH Prepend Database.

81. A packet switched routing control system, comprising:
- a) a plurality of PNAPs for creating route configurations and having PNAP routing packets;
 - b) a network service provider having connections to a said plurality of PNAPs,
5 said PNAP routing packets between a PNAP customer and a destination within said network service provider over said connection to said plurality of PNAPs, said routing being symmetrical over a forward and return routing path.
82. A packet switched routing control method, comprising the steps of:
- a) creating a plurality of PNAPs to facilitate route configurations;
 - b) connecting a network service provider to said plurality of PNAPs; and
 - c) routing packets between a PNAP customer and a destination within said
5 network service provider over said connection to said PNAP, said routing being symmetrical over a forward and return routing path.




Figure 1




Figure 2

3/4




Figure 3

4/4




Figure 4

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US00/22470

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) :H04L 12/46
US CL :370/255, 351, 401

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : Please See Extra Sheet.

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

Easi
search terms: access, packet routing

C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|-----------|--|-----------------------|
| A | US 5,477,536 A (PICARD) 19 DECEMBER 1995, ALL | 1-82 |
| A | US 5,452,294 A (NATARAJAN) 19 SEPTEMBER 1995, ALL | 1-82 |

Further documents are listed in the continuation of Box C. See patent family annex.

| | | |
|---|-----|--|
| • Special categories of cited documents: | "T" | later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention |
| *A* document defining the general state of the art which is not considered to be of particular relevance | "X" | document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone |
| *E* earlier document published on or after the international filing date | "Y" | document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art |
| *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) | "&" | document member of the same patent family |
| *O* document referring to an oral disclosure, use, exhibition or other means | | |
| *P* document published prior to the international filing date but later than the priority date claimed | | |

| | |
|---|---|
| Date of the actual completion of the international search | Date of mailing of the international search report |
| 22 OCTOBER 2000 | 21 NOV 2000 |
| Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231 Facsimile No. (703) 305-3230 | Authorized officer KWANG B. YAO <i>Regenia Zagan</i> Telephone No. (703) 308-0000 |

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US00/22470

B. FIELDS SEARCHED

Minimum documentation searched

Classification System: U.S.

370/248,252,253,254,255,351,400,401,402,403,404,405; 709/220,221,222,223,238,239,240,241,242